

E-Center - Feature # 1939: Enhanced status call to DRS

Status:	Feedback	Priority:	Normal
Author:	David Eads	Category:	Scalable Backend
Created:	09/27/2011	Assigned to:	Maxim Grigoriev
Updated:	09/28/2011	Due date:	
Subject:	Enhanced status call to DRS		
Description:	<p>Max and I have been emailing a bit about an enhanced status call for the DRS, which will be useful in many ways.</p> <p>Here's the current data structure on the table, which I have some concerns with:</p> <pre><pre> { gearman: { scv: { xenmon.fnal.gov: { 10161: { dispatch_data: { running: "0", available: "6", queued: "0" } } } } production: { xenmon.fnal.gov: { 10121: { dispatch_data: { running: "0", ... } } } } } } } } } } } } </pre></pre> <p>My questions/concerns:</p> <ul style="list-style-type: none">* Should the hostname be a key in the data structure at this point? Will there ever be more than one host associated with a "service name"?* Can we include external listening port in the data structure somewhere? That's pretty important from my side of things.* I'm wary of using a "service name" such as "production" or "scv" without a little more rigid definition of what it means and what the allowed values are. I like hostname + external listening port as it is unique and agnostic to the purpose or capabilities of the underlying service.* Can the different running DRS instances be configured to have different functionality? Right now, I can query both the SCV service with a regular data query and the "production" service with a site query. This seems like it could cause some issues, especially given the ease with which the services can be directly queried.		

History

09/27/2011 03:31 pm - Maxim Grigoriev

- Status changed from New to Resolved

hostname is not important and could be anything, the important is what is the role of the DRS. The external port also is not important and can change. Just use scv for one type of queries and production for everything else. And yes, the DRS for SCV will be updated to not support other calls at all. The same is already done for the "monitoring" instance on port 10503. You will get :

```
{ error: "not supported" }
```

09/27/2011 03:45 pm - David Eads

- Status changed from Resolved to Feedback

The hostname and external port could wind up being pretty important for service discovery. It seems like a data structure that makes that info a property and not a high-level key will be easier to parse and use:

```
<pre>
scv: {
  status: "ok",
  service: {
    hostname: "xenmon.fnal.gov",
    port: 8066,
    gearman_port: 10161,
  },
  services: {
    dispatch_data: {
      running: "0",
      available: "6",
      queued: "0"
    }, ...
  }
},
production: {...}
</pre>
```

Also the 'gearman' key seems completely unnecessary -- what's the purpose?

09/28/2011 10:52 am - Maxim Grigoriev

"gearman" key is here because this is just one piece of infrastructure. I am planning to add more monitoring details with other keys - db, server. Single DRS may work with multiple gearman daemons and multiple ports/hosts. Each "port" will have the same list of registered workers, it is required to check all of them for any abnormalities. The DRS port will not be returned because it is controlled by other part - nginx config and known only to the frontend.

09/28/2011 11:15 am - David Eads

Thanks for the explanation. It is a huge bummer about the listening ports, though, in terms of being able to do service discovery and to connect the frontend configuration (which only knows about the public ports) to the backend status report. Without that, the user will have to know quite a bit about their DRS configuration to set up the front-end (which seems perfectly acceptable for now, but could be a longer term issue if there's interest in deploying either component elsewhere).

One minor gripe about service naming: We should come up with a consistent, informative way to name them: "scv" and "production" are a little vague. Something like "drs", "drs_path", "drs_site_centric", and "ads" might be a little better.